APPENDIX C

BRENNAN CENTER TASK FORCE ON VOTING SYSTEM SECURITY, LAWRENCE NORDEN, CHAIR

Excerpted from p.p. 8-19

METHODOLOGY

The Task Force concluded, and the peer review team at NIST agreed, that the best approach for comprehensively evaluating voting system threats was to: (1) identify and categorize the potential threats against voting systems, (2) prioritize these threats based upon an agreed upon metric (which would tell us how difficult each threat is to accomplish from the attacker's point of view), and (3) determine, utilizing the same metric employed to prioritize threats, how much more difficult each of the catalogued attacks would become after various sets of countermeasures are implemented.

This model allows us to identify the attacks we should be most concerned about (*i.e.,* the most practical and least difficult attacks). Furthermore, it allows us to quantify the potential effectiveness of various sets of countermeasures (*i.e.,* how difficult the least difficult attack is after the countermeasure has been implemented). Other potential models considered, but ultimately rejected by the Task Force, are detailed in Appendix B.

IDENTIFICATION OF THREATS

The first step in creating a threat model for voting systems was to identify as many potential attacks as possible. To that end, the Task Force, together with the participating election officials, spent several months identifying voting system vulnerabilities. Following this work, NIST held a Voting Systems Threat Analysis Workshop on October 7, 2005. Members of the public were invited to write up and post additional potential attacks. Taken together, this work produced over 120 potential attacks on the three voting systems. They are detailed in the catalogs annexed.[20] Many of the attacks are described in more detail at http://vote.nist.gov/threats/papers.htm.

The types of threats detailed in the catalogs can be broken down into nine categories: (1) the insertion of corrupt software into machines prior to Election Day; (2) wireless and other remote control attacks on voting machines on Election Day; (3) attacks on tally servers; (4) miscalibration of voting machines; (5) shut off of voting machine features intended to assist voters; (6) denial of service attacks; (7) actions by corrupt poll workers or others at the polling place to affect votes cast; (8) vote buying schemes; (9) attacks on ballots or VVPT. Often, the actual attacks involve some combination of these categories. We provide a discussion of each type of attack in "Categories of Attacks," *infra* at pp. 24–27.

PRIORITIZING THREATS: NUMBER OF INFORMED PARTICIPANTS AS METRIC

Without some form of prioritization, a compilation of the threats is of limited value. Only by prioritizing these various threats could we help election officials identify which attacks they should be most concerned about, and what steps could be taken to make such attacks as difficult as possible. As discussed below, we have determined the level of difficulty for each attack where the attacker is attempting to affect the outcome of a close statewide election.[21]

There is no perfect way to determine which attacks are the least difficult, because each attack requires a different mix of resources – well-placed insiders, money, programming skills, security expertise, *etc*. Different attackers would find certain resources easier to acquire than others. For example, election fraud committed by local election officials would always involve well-placed insiders and a thorough understanding of election procedures; at the same time, there is no reason to expect such officials to have highly skilled hackers or first-rate programmers working with them. By contrast, election fraud carried out by a foreign government would likely start with plenty of money and technically skilled attackers, but probably without many conveniently placed insiders or detailed knowledge of election procedures.

Ultimately, we decided to use the "number of informed participants" as the metric for determining attack difficulty. An attack which uses fewer participants is deemed the easier attack.

We have defined "informed participant" as someone whose participation is needed to make the attack work, and who knows enough about the attack to foil or expose it. This is to be distinguished from a participant who unknowingly assists the attack by performing a task that is integral to the attack's successful execution without understanding that the task is part of an attack on voting systems.

The reason for using the security metric "number of informed participants" is relatively straightforward: the larger a conspiracy is, the more difficult it would be to keep it secret. Where an attacker can carry out an attack by herself, she need only trust herself. On the other hand, a conspiracy that requires thousands of people to take part (like a vote-buying scheme) also requires thousands of people to keep quiet. The larger the number of people involved, the greater the likelihood that one of them (or one who was approached, but declined to take part) would either inform the public or authorities about the attack, or commit some kind of error that causes the attack to fail or become known.

Moreover, recruiting a large number of people who are willing to undermine the integrity of a statewide election is also presumably difficult. It is not hard to imagine two or three people agreeing to work to change the outcome of an election. It seems far less likely that an attacker could identify and employ hundreds or thousands of similarly corrupt people without being discovered.

We can get an idea of how this metric works by looking at one of the threats listed in our catalogs: the vote-buying threat, where an attacker or attackers pay individuals to vote for a particular candidate. This is Attack Number 26 in the PCOS Attack Catalog[22] (though this attack would not be substantially different against DREs or DREs w/ VVPT).[23] In order to work under our current types of voting systems, this attack requires (1) at least one person to purchase votes, (2) many people to agree to sell their votes, and (3) some way for the purchaser to confirm that the voters she pays actually voted for the candidate she supported. Ultimately, we determined that, while practical in smaller contests, a vote-buying attack would be an exceptionally difficult way to affect the outcome of a statewide election. This is because, even in a typically close statewide election, an attacker would need to involve thousands of voters to ensure that she could affect the outcome of a statewide race.[24]

For a discussion of other metrics we considered, but ultimately rejected, *see* Appendix C.

DETERMINING NUMBER OF INFORMED PARTICIPANTS

DETERMINING THE STEPS AND VALUES FOR EACH ATTACK

The Task Force members broke down each of the catalogued attacks into its necessary steps. For instance, Attack 12 in the PCOS Attack Catalog is “Stuffing Ballot Box with Additional Marked Ballots.”[25] We determined that, at a minimum, there were three component parts to this attack: (1) stealing or creating the ballots and then marking them, (2) scanning marked ballots through the PCOS scanners, probably before the polls opened, and (3) modifying the poll books in each location to ensure that the total number of votes in the ballot boxes was not greater than the number of voters who signed in at the polling place.

Task Force members then assigned a value representing the minimum number of persons they believed would be necessary to accomplish each goal. For PCOS Attack 12, the following values were assigned:[26]

Minimum number required to steal or create ballots: 5 persons total.[27]

Minimum number required to scan marked ballots: 1 per polling place attacked.

Minimum number required to modify poll books: 1 per polling place attacked.[28]

After these values were assigned, the Brennan Center interviewed several election officials to see whether they agreed with the steps and values assigned to each attack.[29] When necessary, the values and steps were modified. The new catalogs, including attack steps and values, were then reviewed by Task Force members. The purpose of this review was to ensure, among other things, that the steps and values were sound.

These steps and values tell us how difficult it would be to accomplish a *single attack in a single polling place.* They do not tell us how many people it would take to change the outcome of an election successfully – that depends, of course, on specific facts about the jurisdiction: how many votes are generally recorded in each polling place, how many polling places are there in the jurisdiction, and how close is the race? For this reason, we determined that it was necessary to construct a hypothetical jurisdiction, to which we now turn.

NUMBER OF INFORMED PARTICIPANTS NEEDED TO CHANGE STATEWIDE ELECTION

We have decided to examine the difficulty of each attack in the context of changing the outcome of a reasonably close statewide election. While we are concerned by potential attacks on voting systems in any type of election, we are most troubled by attacks that have the potential to affect large numbers of votes. These are the attacks that could actually change the outcome of a statewide election with just a handful of attack participants.

We are less troubled by attacks on voting systems that can only affect a small number of votes (and might therefore be more useful in local elections). This is because there are many non-system attacks that can also affect a small number of votes (*i.e.*, sending out misleading information about polling places, physically intimidating voters, submitting multiple absentee ballots, *etc.*). Given the fact that these non-system attacks are likely to be less difficult in terms of number of participants, financial cost, risk of detection, and time commitment, we are uncertain that an attacker would target *voting machines* to alter a small number of votes.

In order to evaluate how difficult it would be for an attacker to change the outcome of a statewide election, we created a composite jurisdiction. The composite jurisdiction was created to be representative of a relatively close statewide election. We did not want to examine a statewide election where results were so skewed toward one candidate (for instance, the re-election of Senator Edward M. Kennedy in 2000, where he won 73% of the vote[30]), that reversing the election results would be impossible without causing extreme public suspicion. Nor did we want to look at races where changing only a relative handful of votes (for instance, the Governor's race in Washington State in 2004, which was decided by a mere 129 votes[31]) could affect the outcome of an election; under this scenario, many of the potential attacks would involve few people, and therefore look equally difficult.

We have named our composite jurisdiction "the State of Pennasota." The State of Pennasota is a composite of ten states: Colorado, Florida, Iowa, Ohio, New Mexico, Pennsylvania, Michigan, Nevada, Wisconsin and Minnesota. These states were chosen because they were the ten "battleground" states that Zogby International consistently polled in the spring, summer, and fall 2004.[32] These are statewide elections that an attacker would have expected, ahead of time, to be fairly close.

We have also created a composite election, which we label the "Governor's Race" in Pennasota. The results of this election are a composite of the actual results in the same ten states in the 2004 Presidential Election.

We have used these composites as the framework by which to evaluate the difficulty of the various catalogued attacks.[33] For instance, we know a ballot-box stuffing attack would require roughly five people to create and mark fake ballots, as well as one person per polling place to stuff the boxes, and one person per polling place to modify the poll books. But, in order to determine how many informed participants would be needed to affect a statewide race, we need to know how many polling places would need to be attacked.

The composite jurisdiction and composite election provide us with information needed to answer these questions: *i.e.*, how many extra votes our attackers would need to add to their favored candidate's total for him to win, how many ballots our attackers can stuff into a particular polling place's ballot box without arousing suspicion (and related to this, how many votes are generally cast in the average polling place), how many polling places are there in the state, *etc.* We provide details about both the composite jurisdiction and election in the section entitled "Governor's Race, State of Pennasota, 2007," *infra* at pp 20–27.

LIMITS OF INFORMED PARTICIPANTS AS METRIC

Of the possible metrics we considered, we believe that measuring the number of people who know they are involved in an attack (and thus could provide evidence of the attack to the authorities and/or the media), is the best single measure of attack difficulty; as already discussed, we have concluded that the more people an attacker is forced to involve in his attack, the more likely it is that one of the participants would reveal the attack's existence and foil the attack, perhaps sending attackers to jail. However, we are aware of a number of places where the methodology could provide us with questionable results.

By deciding to concentrate on size of attack team, we mostly ignore the need for other resources when planning an attack. Thus, a software attack on DREs which makes use of steganography[34] to hide attack instruction files (*see* "DRE w/ VVPT Attack No.1a", discussed in greater detail, *infra* at pp. 62–65) is considered easier than an attack program delivered over a wireless network at the polling place (*see* discussion of wireless networks, *infra* at pp. 85–91). However, the former attack probably requires a much more technologically sophisticated attacker.

Another imperfection with this metric is that we do not have an easy way to represent how much choice the attacker has in finding members of his attack team. Thus, with PCOS voting, we conclude that the cost of subverting a routine audit of ballots is roughly equal to the cost of intercepting ballot boxes in transit and substituting altered ballots (*see* discussion of PCOS attacks, *infra* at pp. 77–83). However, subverting the audit team requires getting a specific set of trusted people to cooperate with the attacker. By contrast, the attacker may be able to decide which precincts to tamper with based on which people he has already recruited for his attack.

In an attempt to address this concern, we considered looking at the number of "insiders" necessary to take part in each attack. Under this theory, getting five people to take part in a conspiracy to attack a voting system might not be particularly difficult. But getting five well-placed county election officials to take part in the attack would be (and should be labeled) the more difficult of the two attacks. Because, for the most part, the low-cost attacks we have identified do not necessarily involve well placed insiders (but could, for instance, involve one of many people with access to commercial off the shelf software ("COTS") during development or at the vendor), we do not believe that using this metric would have substantially changed our analysis.[35]

Finally, these attack team sizes do not always capture the logistical complexity of an attack. For example, an attack on VVPT machines involving tampering with the voting machine software and also replacing the paper records in transit requires the attacker to determine what votes were falsely produced by the voting machine and print replacement records in time to substitute them. While this is clearly possible, it raises a lot of operational difficulties – a single failed substitution leaves the possibility that the attack would be detected during the audit of ballots.

We have tried to keep these imperfections in mind when analyzing and discussing our least difficult attacks.

We suspect that much of the disagreement between voting officials and computer security experts in the last several years stems from a difference of opinion in prioritizing the difficulty of attacks. Election officials, with extensive experience in the logistics of handling tons of paper ballots, have little faith in paper and understand the kind of breakdowns in procedures that lead to traditional attacks like ballot box stuffing; in contrast, sophisticated attacks on computer voting systems appear very difficult to many of them. Computer security experts understand sophisticated attacks on computer systems, and recognize the availability of tools and expertise that makes these attacks practical to launch, but have no clear idea how they would manage the logistics of attacking a paper-based system. Looking at attack team size is one way to bridge this difference in perspective.

EFFECTS OF IMPLEMENTING COUNTERMEASURE SETS

The final step of our threat analysis is to measure the effect of certain countermeasures against the catalogued attacks. How much more difficult would the attacks become once the countermeasures are put into effect? How many more informed participants (if any) would be needed to counter or defeat these countermeasures? Our process for examining the effectiveness of a countermeasure mirrors the process for determining the difficulty of an attack: we first asked whether the countermeasure would allow us to detect an attack with near certainty. If we agreed that the countermeasure would expose the attack, we identified the steps that would be necessary to circumvent or defeat the countermeasure. For each step to defeat the countermeasure, we determined the number of additional informed participants (if any) that an attacker would need to add to his team. As with the process for determining attack difficulty, the Brennan Center interviewed numerous election officials to see whether they agreed with the steps and values assigned. When necessary, the values and steps for defeating the countermeasures were altered to reflect the input of election officials.

COUNTERMEASURES EXAMINED

BASIC SET OF COUNTERMEASURES

The first set of countermeasures we looked at is the “Basic Set” of countermeasures. This Basic Set was derived from security survey responses[36] we received from county election officials around the country, as well as additional interviews with more than a dozen current and former election officials. Within the Basic Set of countermeasures are the following procedures:

Inspection

· The jurisdiction is not knowingly using any uncertified software that is subject to inspection by the Independent Testing Authority (often referred to as the "ITA").[37]

Physical Security for Machines

· Ballot boxes (to the extent they exist) are examined (to ensure they are empty) and locked by poll workers immediately before the polls are opened.

· Before and after being brought to the polls for Election Day, voting systems for each county are locked in a single room, in a county warehouse.

· The warehouse has perimeter alarms, secure locks, video surveillance and regular visits by security guards.

· Access to the warehouse is controlled by sign-in, possibly with card keys or similar automatic logging of entry and exit for regular staff.

· Some form of "tamper evident" seals are placed on machines before and after each election.

· The machines are transported to polling locations five to fifteen days before Election Day.

Chain of Custody/Physical Security of Election Day Records

· At close of the polls, vote tallies for each machine are totaled and compared with number of persons that have signed the poll books.

· A copy of totals for each machine is posted at each polling place on Election Night and taken home by poll workers to check against what is posted publicly at election headquarters, on the web, in the papers, or elsewhere.[38]

· All audit information (*i.e.,* Event Logs, VVPT records, paper ballots, machine printouts of totals) that is not electronically transmitted as part of the unofficial upload to the central election office, is delivered in official, sealed and hand-delivered information packets or boxes. All seals are numbered and tamper-evident.

· Transportation of information packets is completed by two election officials representing opposing parties who have been instructed to remain in joint custody of the information packets or boxes from the moment it leaves the precinct to the moment it arrives at the county election center.

· Each polling place sends its information packets or boxes to the county election center separately, rather than having one truck or person pick up this data from multiple polling locations.

· Once the sealed information packets or boxes have reached the county election center, they are logged. Numbers on the seals are checked to ensure that they have not been replaced. Any broken or replaced seals are logged. Intact seals are left intact.

· After the packets and/or boxes have been logged, they are provided with physical security precautions at least as great as those listed for voting machines, above. Specifically, for Pennasota, we have assumed the room in which the packets are stored have perimeter alarms, secure locks, video surveillance and regular visits by security guards and county police officers; and access to the room is controlled by sign-in, possibly with card keys or similar automatic logging of entry and exit for regular staff.

Testing[39]

· An Independent Testing Authority has certified the model of voting machine used in the polling place.

· Acceptance Testing[40] is performed on machines at time, or soon after they are received by County.

· Pre-election Logic and Accuracy[41] testing is performed by the relevant election official.

· Prior to opening the polls, every voting machine and vote tabulation system is checked to see that it is still configured for the correct election, including the correct precinct, ballot style, and other applicable details.

REGIMEN FOR AUTOMATIC ROUTINE AUDIT PLUS BASIC SET OF COUNTERMEASURES.

The second set of countermeasures is the Regimen for an Automatic Routine Audit Plus Basic Set of Countermeasures.

Some form of routine auditing of voter-verified paper records occurs in 12 states, to test the accuracy of electronic voting machines. They generally require between 1 and 10% of all precinct voting machines to be audited after each election. [42]

Jurisdictions can implement this set of countermeasures only if their voting systems produce some sort of voter-verified paper record of each vote. This could be in the form of a paper ballot, in the case of PCOS, or a voter-verified paper trail (“VVPT”), in the case of DREs.

We have assumed that jurisdictions take the following steps when conducting an Automatic Routine Audit (when referring to this set of assumptions “Regimen for an Automatic Routine Audit”):

The Audit

· Leaders of the major parties in each county are responsible for selecting a sufficient number of audit-team members to be used in that county.[43]

· Using a highly transparent random selection mechanism (*see* point ii, below), the voter-verified paper records for between a small percentage of all voting machines in the State are selected for auditing.

· Using a transparent random selection method, auditors are assigned to the selected machines (two or three people, with representatives of each major political party, would comprise each audit team).

· The selection of voting machines, and the assignment of auditors to machines, occurs immediately before the audits take place. The audits take place as soon after polls close as possible – for example, at 9 a.m. the morning after polls close.

· Using a transparent random selection method, county police officers, security personnel and the video monitor assigned to guard the voter-verified records are chosen from a large pool of on-duty officers and employees on election night.

· The auditors are provided the machine tallies and are able to see that the county tally reflects the sums of the machine tallies before the start of the inspection of the paper.

· The audit would include a tally of spoiled ballots (in the case of VVPT, the number of cancellations recorded), overvotes, and undervotes.

Transparent Random Selection Process

In this report, we have assumed that random auditing procedures are in place for both the Regimen for an Automatic Routine Audit and Regimen for Parallel Testing. We have further assumed procedures to prevent a single, corrupt person from being able to fix the results. This implies a kind of transparent and public random procedure.

For the Regimen for an Automatic Routine Audit there are at least two places where transparent, random selection processes are important: in the selection of precincts to audit, and in the assignment of auditors to the precincts they will be auditing.

Good election security can employ Transparent Random Selection in other places with good effect:

· the selection of parallel testers from a pool of qualified individuals.

· the assignment of police and other security professionals from on-duty lists, to monitor key materials, for example, the VVPT records between the time that they arrive at election central and the time of the completion of the ARA.

If a selection process for auditing is to be trustworthy and trusted, ideally:

· The whole process will be publicly observable or videotaped;[44]

· The random selection will be publicly verifiable, *i.e.*, anyone observing will be able to verify that the sample was chosen randomly (or at least that the number selected is not under the control of any small number of people); and

· The process will be simple and practical within the context of current election

practice so as to avoid imposing unnecessary burdens on election officials.
There are a number of ways that election officials can ensure some kind of transparent randomness. One way would be to use a state lottery machine to select precincts or polling places for auditing. We have included two potential examples of transparent random selection processes in Appendix F. These apply to the Regimen for Parallel Testing as well.

REGIMEN FOR PARALLEL TESTING PLUS BASIC SET OF COUNTERMEASURES

The final set of countermeasures we have examined is “Parallel Testing” plus the Basic Set of countermeasures. Parallel Testing, also known as election-day testing, involves selecting voting machines at random and testing them as realistically as possible during the period that votes are being cast.

Parallel Testing

In developing our set of assumptions for Parallel Testing, we relied heavily upon interviews with Jocelyn Whitney, Project Manager for Parallel Testing in the State of California, and conclusions drawn from this Report.[45] In our analysis, we assume that the following procedures would be included in the Parallel Testing regimen (when referring to this regimen “Regimen for Parallel Testing”) that we evaluate:

· At least two of each DRE model (meaning both vendor and model) would be selected for Parallel Testing;

· At least two DREs from each of the three largest counties would be parallel tested;

· Counties to be parallel tested would be chosen by the Secretary of State in a transparent and random manner.

· Counties would be notified as late as possible that machines from one of their precincts would be selected for Parallel Testing;[46]

· Precincts would be selected through a transparent random mechanism;

· A video camera would record testing;

· For each test, there would be one tester and one observer;

· Parallel Testing would occur at the polling place;

· The script for Parallel Testing would be generated in a way that mimics voter behavior and voting patterns for the polling place;

· At the end of the Parallel Testing, the tester and observer would reconcile vote totals in the script with vote totals reported on the machine.

Transparent Random Selection Process

We further assume that the same type of transparent random selection process that would be used for the Regimen for Automatic Routine Audit would also be employed for the Regimen for Parallel Testing to determine which machines would be subjected to testing on Election Day.

APPENDIX C

ALTERNATIVE SECURITY METRICS CONSIDERED

Dollars Spent

The decision to use the number of informed participants as the metric for attack level difficulty came after considering several other potential metrics. One of the first metrics we considered was the dollar cost of attacks. This metric makes sense when looking at attacks that seek financial gain – for instance, misappropriating corporate funds. It is not rational to spend $100,000 on the misappropriation of corporate funds if the total value of those funds is $90,000. Ultimately, we rejected this metric as the basis for our analysis because the dollar cost of the attacks we considered were dwarfed by both (1) current federal and state budgets, and (2) the amounts currently spent legally in state and federal political campaigns.

Time of Attack

The relative security of safes and other safety measures are often rated in terms of "time to defeat." This was rejected as metric of difficulty because it did not seem relevant to voting systems. Attackers breaking into a house are concerned with the amount of time it might take to complete their robbery because the homeowners or police might show up. With regard to election fraud, many attackers may be willing to start months or years before an election if they believe they can control the outcome. As discussed *supra* at pp. 35–48, attackers may be confident that they can circumvent the independent testing authorities and other measures meant to identify attacks, so that the amount of time an attack takes becomes less relevant.